

**Exam Code:** SPLK-3002

**Exam Name:** Splunk IT Service Intelligence Certified Admin



## Exam A

### QUESTION 1

ITSI Saved Search Scheduling is configured to use `realtime_schedule = 0`. Which statement is accurate about this configuration?

- A. If this value is set to 0, the scheduler bases its determination of the next scheduled search execution time on the current time.
- B. If this value is set to 0, the scheduler bases its determination of the next scheduled search on the last search execution time.
- C. If this value is set to 0, the scheduler may skip scheduled execution periods.
- D. If this value is set to 0, the scheduler might skip some execution periods to make sure that the scheduler is executing the searches running over the most recent time range.

**Correct Answer: B**

**Section:**

**Explanation:**

ITSI Saved Search Scheduling is a feature that allows you to schedule searches that run periodically to populate the data for your KPIs. You can configure various settings for your scheduled searches, such as the search frequency, the time range, the cron expression, and so on. One of the settings is `realtime_schedule`, which controls the way the scheduler computes the next execution time of a scheduled search. The statement that is accurate about this configuration is:

B) If this value is set to 0, the scheduler bases its determination of the next scheduled search on the last search execution time. This is called continuous scheduling. If set to 0, the scheduler never skips scheduled execution periods. However, the execution of the saved search might fall behind depending on the scheduler's load. Use continuous scheduling whenever you enable the summary index option.

The other statements are not accurate because:

A) If this value is set to 0, the scheduler bases its determination of the next scheduled search execution time on the current time. This is not true because this is what happens when the value is set to 1, not 0.

C) If this value is set to 0, the scheduler may skip scheduled execution periods. This is not true because this is what happens when the value is set to 1, not 0.

D) If this value is set to 0, the scheduler might skip some execution periods to make sure that the scheduler is executing the searches running over the most recent time range. This is not true because this is what happens when the value is set to 1, not 0.

### QUESTION 2

What effects does the KPI importance weight of 11 have on the overall health score of a service?

- A. At least 10% of the KPIs will go critical.
- B. Importance weight is unused for health scoring.
- C. The service will go critical.
- D. It is a minimum health indicator KPI.

**Correct Answer: B**

**Section:**

**Explanation:**

The KPI importance weight is a value that indicates how much a KPI contributes to the overall health score of a service. The importance weight can range from 1 (lowest) to 10 (highest). The statement that applies when configuring a KPI importance weight of 11 is:

B) Importance weight is unused for health scoring. This is true because an importance weight of 11 is invalid and cannot be used for health scoring. The maximum value for importance weight is 10.

The other statements do not apply because:

A) At least 10% of the KPIs will go critical. This is not true because an importance weight of 11 does not affect the severity level of any KPIs.

C) The service will go critical. This is not true because an importance weight of 11 does not affect the health score or status of any service.

D) It is a minimum health indicator KPI. This is not true because an importance weight of 11 does not indicate anything about the minimum health level of a KPI.

### QUESTION 3

Which of the following is an advantage of using adaptive time thresholds?

- A. Automatically update thresholds daily to manage dynamic changes to KPI values.
- B. Automatically adjust KPI calculation to manage dynamic event data.
- C. Automatically adjust aggregation policy grouping to manage escalating severity.
- D. Automatically adjust correlation search thresholds to adjust sensitivity over time.

**Correct Answer: A**

**Section:**

**Explanation:**

Adaptive thresholds are thresholds calculated by machine learning algorithms that dynamically adapt and change based on the KPI's observed behavior. Adaptive thresholds are useful for monitoring KPIs that have unpredictable or seasonal patterns that are difficult to capture with static thresholds. For example, you might use adaptive thresholds for a KPI that measures web traffic volume, which can vary depending on factors such as holidays, promotions, events, and so on. The advantage of using adaptive thresholds is:

A) Automatically update thresholds daily to manage dynamic changes to KPI values. This is true because adaptive thresholds use historical data from a training window to generate threshold values for each time block in a threshold template. Each night at midnight, ITSI recalculates adaptive threshold values for a KPI by organizing the data from the training window into distinct buckets and then analyzing each bucket separately. This way, the thresholds reflect the most recent changes in the KPI data and account for any anomalies or trends.

The other options are not advantages of using adaptive thresholds because:

B) Automatically adjust KPI calculation to manage dynamic event data. This is not true because adaptive thresholds do not affect the KPI calculation, which is based on the base search and the aggregation method. Adaptive thresholds only affect the threshold values that are used to determine the KPI severity level.

C) Automatically adjust aggregation policy grouping to manage escalating severity. This is not true because adaptive thresholds do not affect the aggregation policy, which is a set of rules that determines how to group notable events into episodes. Adaptive thresholds only affect the threshold values that are used to generate notable events based on KPI severity level.

D) Automatically adjust correlation search thresholds to adjust sensitivity over time. This is not true because adaptive thresholds do not affect the correlation search, which is a search that looks for relationships between data points and generates notable events. Adaptive thresholds only affect the threshold values that are used by KPIs, which can be used as inputs for correlation searches.

**QUESTION 4**

After a notable event has been closed, how long will the meta data for that event remain in the KV Store by default?

- A. 6 months.
- B. 9 months.
- C. 1 year.
- D. 3 months.

**Correct Answer: A**

**Section:**

**Explanation:**

By default, notable event metadata is archived after six months to keep the KV store from growing too large.

**QUESTION 5**

Which of the following is a best practice for identifying the most effective services with which to start an iterative ITSI deployment?

- A. Only include KPIs if they will be used in multiple services.
- B. Analyze the business to determine the most critical services.
- C. Focus on low-level services.
- D. Define a large number of key services early.

**Correct Answer: B**

**Section:**

**Explanation:**

A best practice for identifying the most effective services with which to start an iterative ITSI deployment is to analyze the business to determine the most critical services that have the most impact on revenue, customer satisfaction, or other key performance indicators. You can use the Service Analyzer to prioritize and monitor these services.

Reference:Service Analyzer

**QUESTION 6**

When creating a custom deep dive, what color are services/KPIs in maintenance mode within the topology view?

- A. Gray
- B. Purple
- C. Gear Icon
- D. Blue

**Correct Answer: A**

**Section:**

**Explanation:**

When creating a custom deep dive, services or KPIs that are in maintenance mode are shown in gray color in the topology view. This indicates that they are not actively monitored and do not generate alerts or notable events.

Reference: Deep Dives

**QUESTION 7**

Which deep dive swim lane type does not require writing SPL?

- A. Event lane.
- B. Automatic lane.
- C. Metric lane.
- D. KPI lane.

**Correct Answer: D**

**Section:**

**Explanation:**

A KPI lane is a type of deep dive swim lane that does not require writing SPL. You can simply select a service and a KPI from a drop-down list and ITSI will automatically populate the lane with the corresponding data. You can also adjust the threshold settings and time range for the KPI lane.

Reference: [KPI Lanes]

**QUESTION 8**

Which of the following items apply to anomaly detection? (Choose all that apply.)

- A. Use AD on KPIs that have an unestablished baseline of data points. This allows the ML pattern to perform its magic.
- B. A minimum of 24 hours of data is needed for anomaly detection, and a minimum of 4 entities for cohesive analysis.
- C. Anomaly detection automatically generates notable events when KPI data diverges from the pattern.
- D. There are 3 types of anomaly detection supported in ITSI: adhoc, trending, and cohesive.

**Correct Answer: B, C**

**Section:**

**Explanation:**

Anomaly detection is a feature of ITSI that uses machine learning to detect when KPI data deviates from a normal pattern. The following items apply to anomaly detection:

- B) A minimum of 24 hours of data is needed for anomaly detection, and a minimum of 4 entities for cohesive analysis. This ensures that there is enough data to establish a baseline pattern and compare different entities within a service.
- C) Anomaly detection automatically generates notable events when KPI data diverges from the pattern. You can configure the sensitivity and severity of the anomaly detection alerts and assign them to episodes or teams.

Reference: [Anomaly Detection]

**QUESTION 9**

Which of the following is a best practice when configuring maintenance windows?

- A. Disable any glass tables that reference a KPI that is part of an open maintenance window.
- B. Develop a strategy for configuring a service's notable event generation when the service's maintenance window is open.
- C. Give the maintenance window a buffer, for example, 15 minutes before and after actual maintenance work.
- D. Change the color of services and entities that are part of an open maintenance window in the service analyzer.

**Correct Answer: C**

**Section:**

**Explanation:**

It's a best practice to schedule maintenance windows with a 15- to 30-minute time buffer before and after you start and stop your maintenance work.

A maintenance window is a period of time when a service or entity is undergoing maintenance operations or does not require active monitoring. It is a best practice to schedule maintenance windows with a 15- to 30-minute time buffer before and after you start and stop your maintenance work. This gives the system an opportunity to catch up with the maintenance state and reduces the chances of ITSI generating false positives during maintenance operations. For example, if a server will be shut down for maintenance at 1:00PM and restarted at 5:00PM, the ideal maintenance window is 12:30PM to 5:30PM. The 15- to 30-

minute time buffer is a rough estimate based on 15 minutes being the time period over which most KPIs are configured to search data and identify alert triggers.

Reference:Overview of maintenance windows in ITSI

#### QUESTION 10

In Episode Review, what is the result of clicking an episode's Acknowledge button?

- A. Assign the current user as owner.
- B. Change status from New to Acknowledged.
- C. Change status from New to In Progress and assign the current user as owner.
- D. Change status from New to Acknowledged and assign the current user as owner.

**Correct Answer: D**

**Section:**

**Explanation:**

When an episode warrants investigation, the analyst acknowledges the episode, which moves the status from New to In Progress.

An episode represents a disruption of service operation causing impact to business operations. It is a deduplicated group of notable events occurring as part of a larger sequence, or an incident or period considered in isolation. In Episode Review, you can manage the episodes and their statuses using various actions. One of the actions is Acknowledge, which changes the status of an episode from New to Acknowledged and assigns the current user as the owner. This action indicates that someone is working on resolving the episode and prevents duplicate efforts from other users.

Reference:Overview of Episode Review in ITSI, [Episode actions in Episode Review]

#### QUESTION 11

Which glass table feature can be used to toggle displaying KPI values from more than one service on a single widget?

- A. Service templates.
- B. Service dependencies.
- C. Ad-hoc search.
- D. Service swapping.

**Correct Answer: D**

**Section:**

**Explanation:**

A glass table is a visualization tool that allows you to monitor the interrelationships and dependencies across your IT and business services. You can add metrics like KPIs, ad hoc searches, and service health scores that update in real time against a background that you design. One of the features of glass tables is service swapping, which enables you to toggle displaying KPI values from more than one service on a single widget. You can use service swapping to compare metrics across different services without creating multiple glass tables or widgets.

Reference:Overview of the glass table editor in ITSI, [Configure service swapping on glass tables]

#### QUESTION 12

Which of the following is a characteristic of base searches?

- A. Search expression, entity splitting rules, and thresholds are configured at the base search level.
- B. It is possible to filter to entities assigned to the service for calculating the metrics for the service's KPIs.
- C. The fewer KPIs that share a common base search, the more efficiency a base search provides, and anomaly detection is more efficient.
- D. The base search will execute whether or not a KPI needs it.

**Correct Answer: B**

**Section:**

**Explanation:**

A base search is a search definition that can be shared across multiple KPIs that use the same data source. Base searches can improve search performance and reduce search load by consolidating multiple similar KPIs. One of the characteristics of base searches is that it is possible to filter to entities assigned to the service for calculating the metrics for the service's KPIs. This means that you can use entity filtering rules to specify which entities are relevant for each KPI based on the base search results.

Reference:Create KPI base searches in ITSI, [Filter entities for KPIs based on base searches]

#### QUESTION 13

What are valid ITSI Glass Table editor capabilities? (Choose all that apply.)



- A. Creating glass tables.
- B. Correlation search creation.
- C. Service swapping configuration.
- D. Adding KPI metric lanes to glass tables.

**Correct Answer: A, C, D**

**Section:**

**Explanation:**

Create a glass table to visualize and monitor the interrelationships and dependencies across your IT and business services.

The service swapping settings are saved and apply the next time you open the glass table.

You can add metrics like KPIs, ad hoc searches, and service health scores that update in real time against a background that you design. Glass tables show real-time data generated by KPIs and services.

The glass table editor is a tool that allows you to create and edit glass tables in ITSI. Some of the capabilities of the glass table editor are:

Creating glass tables from scratch or from existing templates.

Configuring service swapping on widgets to toggle displaying metrics from different services.

Adding KPI metric lanes to glass tables to show historical trends of KPI values.

The glass table editor does not support correlation search creation, which is a separate feature in ITSI that allows you to create searches that look for relationships between data points and generate notable events.

Reference: [Overview of the glass table editor in ITSI](#), [\[Configure service swapping on glass tables\]](#), [\[Add KPI metric lanes to glass tables\]](#), [\[Overview of correlation searches in ITSI\]](#)

#### QUESTION 14

Which of the following applies when configuring time policies for KPI thresholds?

- A. A person can only configure 24 policies, one for each hour of the day.
- B. They are great if you expect normal behavior at 1:00 to be different than normal behavior at 5:00
- C. If a person expects a KPI to change significantly through a cycle on a daily basis, don't use it.
- D. It is possible for multiple time policies to overlap.

**Correct Answer: B**

**Section:**

**Explanation:**

Time policies are user-defined threshold values to be used at different times of the day or week to account for changing KPI workloads. Time policies accommodate normal variations in usage across your services and improve the accuracy of KPI and service health scores. For example, if your organization's peak activity is during the standard work week, you might create a KPI threshold time policy that accounts for higher levels of usage during work hours, and lower levels of usage during off-hours and weekends. The statement that applies when configuring time policies for KPI thresholds is:

B) They are great if you expect normal behavior at 1:00 to be different than normal behavior at 5:00. This is true because time policies allow you to define different threshold values for different time blocks, such as AM/PM, work hours/off hours, weekdays/weekends, and so on. This way, you can account for the expected variations in your KPI data based on the time of day or week.

The other statements do not apply because:

- A) A person can only configure 24 policies, one for each hour of the day. This is not true because you can configure more than 24 policies using different time block combinations, such as 3 hour block, 2 hour block, 1 hour block, and so on.
- C) If a person expects a KPI to change significantly through a cycle on a daily basis, don't use it. This is not true because time policies are designed to handle KPIs that change significantly through a cycle on a daily basis, such as web traffic volume or CPU load percent.
- D) It is possible for multiple time policies to overlap. This is not true because you can only have one active time policy at any given time. When you create a new time policy, the previous time policy is overwritten and cannot be recovered.

#### QUESTION 15

What is the main purpose of the service analyzer?

- A. Display a list of All Services and Entities.
- B. Trigger external alerts based on threshold violations.
- C. Allow Analysts to add comments to Alerts.
- D. Monitor overall Service and KPI status.

**Correct Answer: D**

**Section:**

**Explanation:**

The service analyzer is a dashboard that allows you to monitor the overall service and KPI status in ITSI. The service analyzer displays a list of all services and their health scores, which indicate how well each service is performing based on its KPIs. You can also view the status and values of each KPI within a service, as well as drill down into deep dives or glass tables for further analysis. The service analyzer helps you identify issues affecting your services and prioritize them based on their impact and urgency. The main purpose of the service analyzer is:

D) Monitor overall service and KPI status. This is true because the service analyzer provides a comprehensive view of the health and performance of your services and KPIs in real time.

The other options are not the main purpose of the service analyzer because:

A) Display a list of all services and entities. This is not true because the service analyzer does not display entities, which are IT components that require management to deliver an IT service. Entities are displayed in other dashboards, such as entity management or entity health overview.

B) Trigger external alerts based on threshold violations. This is not true because the service analyzer does not trigger alerts, which are notifications sent to external systems or users when certain conditions are met. Alerts are triggered by correlation searches or alert actions configured in ITSI.

C) Allow analysts to add comments to alerts. This is not true because the service analyzer does not allow analysts to add comments to alerts, which are notifications sent to external systems or users

**QUESTION 16**

What is the default importance value for dependent services' health scores?

- A. 11
- B. 1
- C. Unassigned
- D. 10

**Correct Answer: D**

**Section:**

**Explanation:**

By default, impacting service health scores have an importance value of 11.

A service template is a predefined set of KPIs and entity rules that you can apply to a service or a group of services. A service template helps you standardize the configuration and monitoring of similar services across your IT environment. A service template can also include dependent services, which are services that are required for another service to function properly. For example, a web server service might depend on a database service and a network service. The default importance value for dependent services' health scores is:

D) 10. This is true because the importance value indicates how much a dependent service contributes to the health score of the parent service. The default value is 10, which means that the dependent service has the highest impact on the parent service's health score. You can change the importance value of a dependent service in the service template settings.

The other options are not correct because:

A) 11. This is not true because 11 is an invalid value for importance. The valid range is from 1 (lowest) to 10 (highest).

B) 1. This is not true because 1 is the lowest value for importance, not the default value. A value of 1 means that the dependent service has the lowest impact on the parent service's health score.

C) Unassigned. This is not true because every dependent service has an assigned importance value, which defaults to 10.

**QUESTION 17**

What should be considered when onboarding data into a Splunk index, assuming that ITSI will need to use this data?

- A. Use | stats functions in custom fields to prepare the data for KPI calculations.
- B. Check if the data could leverage pre-built KPIs from modules, then use the correct TA to onboard the data.
- C. Make sure that all fields conform to CIM, then use the corresponding module to import related services.
- D. Plan to build as many data models as possible for ITSI to leverage

**Correct Answer: B**

**Section:**

**Explanation:**

When onboarding data into a Splunk index, assuming that ITSI will need to use this data, you should consider the following:

B) Check if the data could leverage pre-built KPIs from modules, then use the correct TA to onboard the data. This is true because modules are pre-packaged sets of services, KPIs, and dashboards that are designed for specific types of data sources, such as operating systems, databases, web servers, and so on. Modules help you quickly set up and monitor your IT services using best practices and industry standards. To use modules, you need to install and configure the correct technical add-ons (TAs) that extract and normalize the data fields required by the modules.

The other options are not things you should consider because:

A) Use | stats functions in custom fields to prepare the data for KPI calculations. This is not true because using | stats functions in custom fields can cause performance issues and inaccurate results when calculating



KPIs. You should use | stats functions only in base searches or ad hoc searches, not in custom fields.

C) Make sure that all fields conform to CIM, then use the corresponding module to import related services. This is not true because not all modules require CIM-compliant data sources. Some modules have their own data models and field extractions that are specific to their data sources. You should check the documentation of each module to see what data requirements and dependencies they have.

D) Plan to build as many data models as possible for ITSI to leverage. This is not true because building too many data models can cause performance issues and resource consumption in your Splunk environment. You should only build data models that are necessary and relevant for your ITSI use cases.

#### QUESTION 18

When changing a service template, which of the following will be added to linked services by default?

- A. Thresholds.
- B. Entity Rules.
- C. New KPIs.
- D. Health score.

**Correct Answer: C**

**Section:**

**Explanation:**

C) New KPIs. This is true because when you add new KPIs to a service template, they will be automatically added to all the services that are linked to that template. This helps you keep your services consistent and up-to-date with the latest KPI definitions.

The other options will not be added to linked services by default because:

A) Thresholds. This is not true because when you change thresholds in a service template, they will not affect the existing thresholds in the linked services. You need to manually apply the threshold changes to each linked service if you want them to inherit the new thresholds from the template.

B) Entity rules. This is not true because when you change entity rules in a service template, they will not affect the existing entity rules in the linked services. You need to manually apply the entity rule changes to each linked service if you want them to inherit the new entity rules from the template.

D) Health score. This is not true because when you change health score settings in a service template, they will not affect the existing health score settings in the linked services. You need to manually apply the health score changes to each linked service if you want them to inherit the new health score settings from the template.

#### QUESTION 19

Which of the following items describe ITSI Deep Dive capabilities? (Choose all that apply.)

- A. Comparing a service's notable events over a time period.
- B. Visualizing one or more Service KPIs values by time.
- C. Examining and comparing alert levels for KPIs in a service over time.
- D. Comparing swim lane values for a slice of time.

**Correct Answer: B, C, D**

**Section:**

**Explanation:**

A deep dive is a dashboard that allows you to analyze the historical trends and anomalies of your KPIs and metrics in ITSI. A deep dive displays a timeline of events and swim lanes of data that you can customize and filter to investigate issues and perform root cause analysis. Some of the capabilities of deep dives are:

B) Visualizing one or more service KPIs values by time. This is true because you can add KPI swim lanes to a deep dive to show the values and severity levels of one or more KPIs over time. You can also compare KPIs from different services or entities using service swapping or entity splitting.

C) Examining and comparing alert levels for KPIs in a service over time. This is true because you can add alert swim lanes to a deep dive to show the alert levels and counts for one or more KPIs over time. You can also drill down into the alert details and view the notable events associated with each alert.

D) Comparing swim lane values for a slice of time. This is true because you can use the time range selector to zoom in or out of a specific time range in a deep dive. You can also use the time brush to select a slice of time and compare the swim lane values for that time period.

The other option is not a capability of deep dives because:

A) Comparing a service's notable events over a time period. This is not true because deep dives do not display notable events, which are alerts generated by ITSI based on certain conditions or correlations. Notable events are displayed in other dashboards, such as episode review or glass tables.

#### QUESTION 20

What is an episode?



- A. A workflow task.
- B. A deep dive.
- C. A notable event group.
- D. A notable event.

**Correct Answer: C**

**Section:**

**Explanation:**

It's a deduplicated group of notable events occurring as part of a larger sequence, or an incident or period considered in isolation.

An episode is a deduplicated group of notable events occurring as part of a larger sequence, or an incident or period considered in isolation. An episode helps you reduce alert noise and focus on the most important issues affecting your IT services. An episode is created by an aggregation policy, which is a set of rules that determines how to group notable events based on certain criteria, such as severity, source, title, and so on. You can use episode review to view, manage, and resolve episodes in ITSI. The statement that defines an episode is:

C) A notable event group. This is true because an episode is composed of one or more notable events that are related by some common factor.

The other options are not definitions of an episode because:

A) A workflow task. This is not true because a workflow task is an action that you can perform on an episode, such as assigning an owner, changing the status, adding comments, and so on.

B) A deep dive. This is not true because a deep dive is a dashboard that allows you to analyze the historical trends and anomalies of your KPIs and metrics in ITSI.

D) A notable event. This is not true because a notable event is an alert generated by ITSI based on certain conditions or correlations, not a group of alerts.

#### QUESTION 21

Which index will contain useful error messages when troubleshooting ITSI issues?

- A. \_introspection
- B. \_internal
- C. itsi\_summary
- D. itsi\_notable\_audit

**Correct Answer: B**

**Section:**

**Explanation:**

The index that will contain useful error messages when troubleshooting ITSI issues is:

B) \_internal. This is true because the \_internal index contains logs and metrics generated by Splunk processes, such as splunkd and metrics.log. These logs can help you diagnose problems with your Splunk environment, including ITSI components and features.

The other indexes will not contain useful error messages because:

A) \_introspection. This is not true because the \_introspection index contains data about Splunk resource usage, such as CPU, memory, disk space, and so on. These data can help you monitor the performance and health of your Splunk environment, but not the error messages.

C) itsi\_summary. This is not true because the itsi\_summary index contains summarized data for your KPIs and services, such as health scores, severity levels, threshold values, and so on. These data can help you analyze the trends and anomalies of your IT services, but not the error messages.

D) itsi\_notable\_audit. This is not true because the itsi\_notable\_audit index contains audit data for your notable events and episodes, such as creation time, owner

